Autonomous Onboard Traverse Science System

Rebecca Castaño, Michele Judd, Tara Estlin, Robert C. Anderson, Lucas Scharenbroich, Lin Song, Daniel Gaines, Forest Fisher, Dominic Mazzoni, Andres Castaño

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, CA 91109
(818) 393-5344
Rebecca.Castaño@jpl.nasa.gov

Abstract—The Onboard Autonomous Science Investigation System (OASIS) is a technology for increasing science return during rover traverses by prioritizing data onboard, and identifying and reacting to unanticipated science opportunities.

Rovers of the future will have the capacity to collect more data than can be downlinked back to Earth. OASIS can increase mission science return by carefully selecting the data with the highest science interest for downlink.

These rovers may also be required to traverse long distances with little to no interaction with the science team on Earth. OASIS can act as a geologist's assistant and can autonomously direct the rover to take additional measurements of "interesting" rocks. The importance of characterizing the terrain along these traverses, a study that is now becoming known as traverse science, increases with the distances the rover must travel.

This paper provides a brief overview of the entire OASIS system and how it analyzes one type of data - grayscale images taken by the rover for engineering and hazard avoidance purposes. Although the OASIS system can apply the same type of analysis to different data types, such as color images, hyperspectral images or point spectrometer data, we will only focus on grayscale images here.

The paper also describes the latest advances in two key aspects of the system: image prioritization and the science alert. In image prioritization, we combine the results from three distinct prioritization methods to arrive at an overall downlink ranking of the images collected during a traverse.

The science alert is a capability that enables the rover to identify and react to a pre-specified, and scientifically significant, signature. Once this signature has been detected via the onboard science analysis component, the planning and scheduling module updates the rover command sequence to stop the traverse and signal Earth of the find. If there is sufficient time and onboard resources before the next downlink opportunity, additional data samples of the target may be autonomously collected.

TABLE OF CONTENTS

1. INTRODUCTION

Rovers offer scientists the ability to move around a planetary surface and explore different areas of interest. The farther the rover can travel, the greater the opportunity exists for increased scientific discovery. In order to reach respectable travel distances, engineers must be able to direct the rover forward in a much more autonomous fashion – without constant "stop, look and explore" directives from the ground based science community.

This "stop, look and explore" phenomenon (resulting from the long two-way communication travel time) was witnessed in the Mars Pathfinder mission during the coordination of the small Sojourner rover. Sojourner traveled approximately 100 meters during its 90-day lifetime [1], but traversed a maximum radial distance of only 12 meters from the lander according to Matt Golombek, Project Scientist for the Mars Pathfinder Mission. Arguably one of the most successful and historic space exploration missions to date, one cannot help but

IEEEAC paper # 1375, Version 5, Updated December 22, 2003.

ponder what other discoveries lay just beyond the 12-meter mark (see Figure 1).

One possibility of extending the rover's mobility exists in a

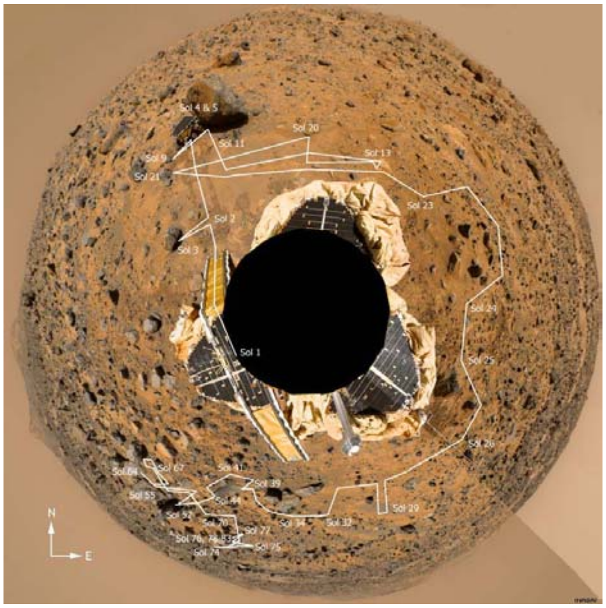

Figure 1 The rover Sojourner's path throughout its 1997 Mars mission.

mission concept of focusing the scientific investigation on multiple science sites that are located several rover traverse days away from each other. In this scenario, scientists would have control over the operations at each of the science sites, and engineers would be given free reign to drive the rover, without interruption, to each of the science sites.

This scenario allows for a deep and concentrated exploration of each science site, but this capability does not come without a price. In order to expeditiously drive the rover to each of the far flung sites, scientists can no longer perform detailed examinations of the long traverses *between* each science site. There is, however, a technology under development to *autonomously* study the terrain during long rover traverses. The Onboard Autonomous Science Investigation System (OASIS) [2,3], a system that is the result of research started at JPL in 1998, was designed to address this type of geological research application. The study of terrain during a traverse is now becoming known as traverse science.

OASIS maximizes the geological data that rovers can collect and subsequently downlink in two key ways: by prioritizing and summarizing all of the data so that the most important data is sent first, and by autonomously searching for pre-specified targets that will trigger additional, opportunistic science measurements.

Prioritization Offsets Downlink Constraints

Due to ongoing technology advancements, the data capture rate of spacecraft instruments is increasing at a rapid pace. Soon spacecraft will be able to collect more data than can be received by NASA's communication-relay antenna infrastructure, the Deep Space Network (DSN). Not only must the DSN maintain the communications from all the current spacecraft in the Solar System, including spacecraft as old as the 26-year-old Voyager spacecraft, but the DSN must also manage the many future missions planned.

Future rovers, therefore, will soon have the capacity to collect more data than can be downlinked back to Earth. Restricting the rover's onboard data collection to accommodate the limited DSN bandwidth of the mission may be the current modus operandi, but it may not be the best option for optimizing the use of the precious bandwidth resource.

OASIS was designed to provide scientists and mission operations personnel with a method of maximizing the science data returned per transmission cycle. With OASIS onboard, the rover is now instructed to collect as much data as it can – regardless of how much of the data the rover can actually downlink back to Earth. OASIS continually reviews and prioritizes the data as it is collected. At the time the data must be sent back to Earth, the most interesting data in the prioritization queue is sent first. As the rover has more data than can be downlinked, OASIS provides the scientists with the most interesting data collected that day.

We have previously presented three data-prioritization techniques, which look for target signatures, novel objects, and objects that are highly representative of the terrain. This paper discusses how the results of these techniques can be efficiently combined into one unified prioritization.

Prioritization of the data gathered during the traverse is important, but so is a short, tabular summary of the entire traverse. OASIS provides a comprehensive overview of all of the rocks that it found during the traverse. This type of information is important to scientists interested in rock type distribution and detection of geologic boundaries. Another advantage of this summary table is that it takes up a very small portion of the available bandwidth, and does not significantly impact the downlink queue. As the summarization component of OASIS was developed previously, it will not be discussed further.

Opportunistic Science – the Science Alert

Besides the obvious advantages of data prioritization and summarization, OASIS has made another contribution to traverse science and it is the science alert.

The science alert is an algorithm designed to find key targets that have been pre-specified by the scientists as important. As data from the rover are fed to the OASIS system, the science alert analyzes the data and looks for a rock that a scientist would want another measurement of, if he/she were actually there on Mars. The characteristics of this rock are programmed into the OASIS system based on the specifications provided by the science team.

If the science alert actually finds a rock that meets these characteristics, it triggers the planning and scheduling system to either stop and call home (an extremely important rock has been found) or schedules another science measurement. Additional science measurements may include a color image of the rock, a spectrometer measurement, or a contact measurement.

This paper addresses the planning and scheduling aspects of the science alert. However, before discussing the details of the development work completed on both the science alert and the unified prioritization technique, a brief overview of OASIS is presented.

2. OASIS SYSTEM OVERVIEW

To assess and subsequently prioritize the scientific value of a set of collected grayscale images, we must first extract the information found within the images. A geologist in the field gets information about a site by identifying geologic features including the albedo, texture, shape, size, color, and arrangement of rocks, and features of the topography such as layers in a cliff face. The geologist analyzes and assesses this data, and then takes some action based on the analysis, such as taking a sample or taking some additional measurement of an interesting rock.

In order for scientists to allow an autonomous system to help investigate the traversed region, the system must be able to perform, albeit in a very simple way, these same types of functions. This system thus acts as a geologist's assistant who helps point out rocks of interest to the geologist.

A schematic of the OASIS system is shown in Figure 2. The color-coded boxes in the figure reveal the three major components that comprise OASIS:

- ***Extract Features from Images:*** Enables extraction of features of interest from collected images of the surrounding terrain. This module both locates rocks in the images and extracts rock properties (features) including shape, texture and albedo.
- ***Analyze and Prioritize Data:*** Uses the extracted features to assess the scientific value of the planetary scene and to generate new science objectives that will further contribute to this assessment. This module consists of three separate prioritization algorithms that analyze the collected data and prioritize the rocks. The results from these three algorithms are then fed into a unified prioritization algorithm that provides two downlink products: a prioritized list of images for downlink and a table that summarizes all of the rocks found on the traverse. A new set of observation goals is also generated to gather further data on rocks that either conform to the pre-set specifications by the science team, or are so novel in comparison to the other rocks, that another data measurement may be required.
- ***Plan and Schedule New Command Sequence:*** Enables dynamic modification of the current rover command sequence (or plan) to accommodate new science requests from the data analysis and

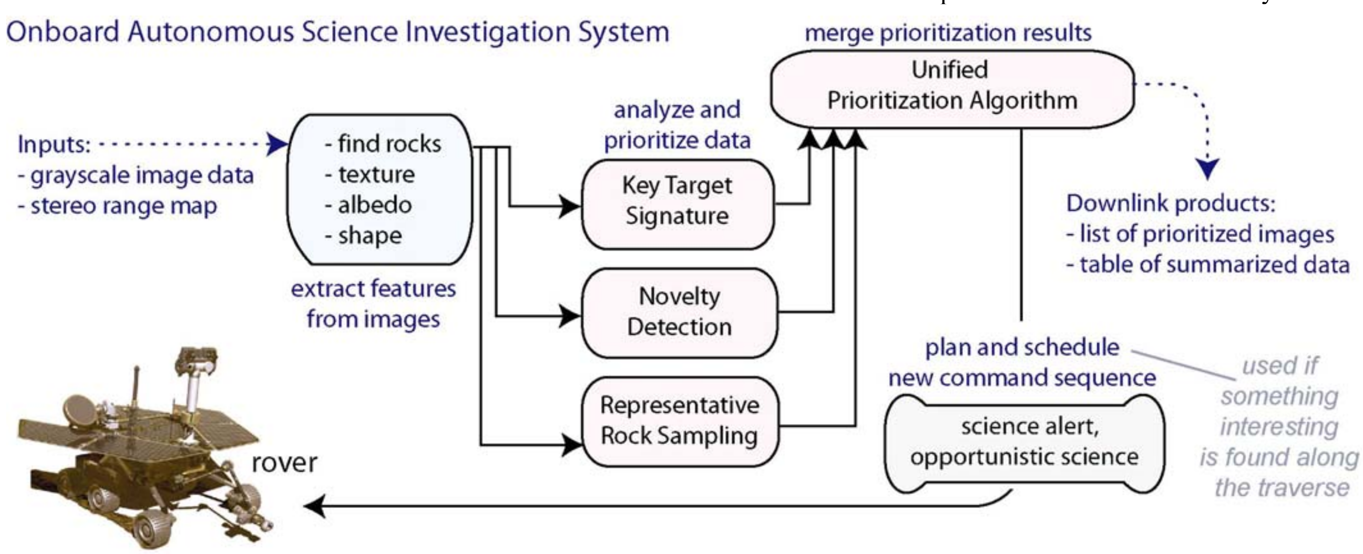

Figure 2 Overview of Onboard Autonomous Science Investigation System (OASIS).

prioritization module. A continuous planning approach is used to iteratively adjust the plan as new goals occur, while ensuring that resource and other operation constraints are met.

3. FEATURE EXTRACTION

As the geologist's assistant, OASIS must be able to find and evaluate the rocks on the surface of Mars.

The first step in the process is to take grayscale image data and a stereo range map and find, and then closely outline, the rocks within the image. JPL has had some success [4] in this area, but as this method required substantial parameter adjustment for each new data set, this was unacceptable for an onboard application. This task is non-trivial (see Figure 3) and the new rock-finding algorithm is still under development.

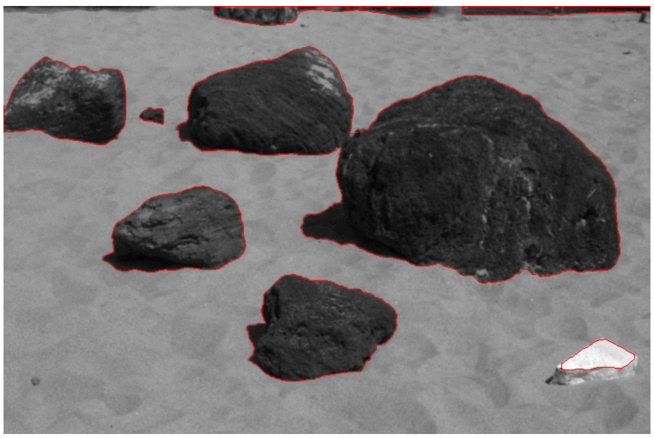

Figure 3 Grayscale image with rocks outlined (in red) by the newer, more robust, version of the rock-finder. Note that the rock-finding algorithm shows the shadows as part of the rocks and that the white rock has not been outlined in its entirety. OASIS is already funded to address these issues.

After OASIS has found the rocks, OASIS extracts the features (or properties) of each rock from the image data. The rock features that are currently extracted include albedo, visual texture and shape. Our scientist collaborators helped us select these rock properties so that what OASIS measures will mirror the properties that an expert would use to evaluate the scientific merit of an image's contents.

We measure albedo, an indicator of the reflectance properties of a surface, by computing the average gray-scale value of the pixels that comprise the image of the rock. The reflectance properties of a rock provide information about its mineralogical composition.

The second rock property extracted is visual texture. Visual texture can provide valuable clues to both the mineral composition and geological history of a rock. Visual texture can be described by gray-scale intensity variations at different orientations and spatial frequencies within the image. We measure texture using a bank of Gabor filters [5,6]. Gabor filters are scale and orientation specific, thus the results of convolving an image with these filters can be successfully used to discriminate between different textures.

Another important and geologically useful feature of rocks is their inherent shape. For example, a rock that is highly rounded may have undergone fluvial processing and traveled far from its source. Conversely, a rock that is highly angular is likely to be close to its source and to have undergone minimal secondary processing. We begin by fitting an ellipse to the boundary points of the identified rock in the image [7]. Our first shape measure is the eccentricity of this ellipse. Our second measure is the error between the boundary points and the ellipse. The third and final measure is angularity, which is measured as the standard deviation of the angle of the edge at each boundary point.

4. DATA PRIORITIZATION

Images with interesting features should be ranked higher than images without distinctive features. We have developed three different prioritization methods that use the extracted rock features to rank the rocks in terms of scientific importance. These three algorithms capture several aspects of science investigation including identification of pre-specified targets, discovery of novel targets and gaining a representative understanding of the data.

The first technique recognizes pre-specified target signatures that have been identified by the science team as data of high interest. This prioritization method enables scientists to efficiently and easily stipulate the value and importance to assign to each feature. Rocks are then prioritized as a function of the distance of their extracted feature vector from the specified weighted feature vector. Scientists are given two ways to set the target signatures that will determine how the rocks are ranked. In the first method, the scientist can directly set the importance of specific feature values. For example, the scientist may chose to prioritize rocks based on two aspects of their shape, such as eccentricity and ellipse fit. The second manner in which scientists can specify a target signature is by selecting a rock with interesting properties from the set of already identified rocks. Rocks that resemble this particular rock in the selected properties are given a high priority.

The second technique, novelty detection, identifies unusual signatures that do not conform to the statistical norm for the region. We have developed three methods for detecting and prioritizing novel rocks, representing the three dominant flavors of machine learning approaches to novelty detection:

- distance-based,
- probability-based (i.e. "generative"),

- and discriminative.

The first novelty detection method is a distance-based k-means clustering approach. Initially, all available rock data is clustered into a specified number (k) of classes. The novelty of any rock is then the distance of the rock feature vector to the nearest center of any of the k clusters. The greater the rock's distance is to the nearest center, the higher the novelty ranking assigned to the rock.

The second technique is a probability-based Gaussian mixture model, which attempts to model the probability density over the feature space. In this approach, the novelty of a rock is inversely proportional to the resulting probability of that rock being generated by the model learned on previous rock data.

The final method is a discrimination-based kernel one-class classifier approach. Here we treat all previous rock data as the "positive class" and learn the discriminant boundary that encloses all that data in the feature space. We essentially consider the previous rock data as a cloud scatter in some D-dimensional space, where D is the number of features. The algorithm learns the boundary of that cloud, so that future rock data that falls farther outside the cloud boundary is considered more novel.

The last prioritization algorithm, known as representative sampling, prioritizes data for downlink by ensuring that representative rocks of the traversed region are returned. One of the objectives for rover traverse science is to gain an understanding of the region being traversed. To meet this objective, the downlink back to Earth should include information on rocks that are typical for a region, and not just information on interesting and unusual rocks. A region is likely populated by several types of rocks with each rock type having a different abundance. If uniform sampling is employed for downlink image selection, as opposed to our autonomous onboard selection process, the downlinked set will be biased towards the dominant class of rock present. This situation may result in smaller classes not being represented at all in the downlinked data.

To provide an understanding of the typical characteristics of a region, rocks are first clustered into groups with similar properties. The data is then prioritized to ensure that representative rocks from each class are sampled. The rocks are clustered into groups based on the features extracted from the image data for each rock. To determine the classes, the property values are concatenated together to form a feature vector, and a weight is assigned to the importance of each property. Different weight assignments can be used as a function of the particular properties that are of interest. For example, albedo and texture are typically used to distinguish types of rocks, but rock size may be used if sorting is of interest. Unsupervised clustering is then used to separate the feature vectors into similar classes. We currently employ k-means due to its relatively low computational requirements, although any unsupervised method could be used. For each class of rocks, we find the most representative rock in the class, i.e., the single rock in any image that is closest to the mean of the set. We give a high priority to the image containing this rock. The optimal number of classes can be determined using cross-validation techniques [8].

The results from all three prioritization techniques must now be merged into one final, ranked list of images for downlink. This is accomplished by the unified prioritization algorithm, which has been developed this year and will be discussed in detail in the next section.

5. UNIFIED PRIORITIZATION ALGORITHM

Three algorithms for prioritizing data based on distinct criteria of science investigation were described. The prioritization information from across these disparate criteria must be combined to determine a downlink priority ranking for the overall data set. Here we describe how we combine the information from the three prioritizations to produce a unified prioritization. Our method accounts for the rankings for a particular criteria as well as the relative weighting of each criteria.

The three algorithms represent methods for evaluating data based on three classes of criteria. For each criterion the data are sorted by how closely they align with the given criteria. For example, in the case of target signature this is the distance of the feature vector to the target feature. Thus, for our three criteria, each data has three fitness scores associated with it indicating its fitness for each of the distinct criteria. The relationship of these values provides more than just ranking information, but also a measure of similarity or significance; however comparisons cannot be made between scores measured based on different criteria, such as novelty vs. target signature. While a combined prioritaztion could be as simple as determining the ranking based on the mean rank of the three algorithms for each data point, our method also consider the fitness scores for the data.

Since fitness scores are not on the same scales across the algorithms, we normalize the fitness values associated with each criterion for all of the data. Further, the ranking or value does not correspond to a linear weighting. Generally, there will be a few examples with high scores and these are the most important. The high scores will fall off in an exponential manner to a plateau where the majority of typical examples reside. These examples are not bad, per se, but are not particularly interesting. Following this plateau will be a smooth drop off to the extremely bad examples which score near zero. The most important object may be significantly more important than the second most important and objects in the lower half may be of virtually no value. We have developed a method with an adjustable non-linear normalization function that implements this. With the adjusted list, the information can then be compared.

Using the normalized fitness scores, the simplest method for prioritizing the data is to rank based on the norm of the fitness vector that has a component for each of the criteria. This method will ensure that data with overall high fitness across the criteria will be given a high priority. It is possible, however, that a number of data will fit better single criterion than any data to the other criteria. This can result in a high priority data set that is biased towards this single criterion and does not contain the best fitting data to other criteria. To ensure that data with high priority in each of the criteria are considered, a procedure for prioritizing that also considers the diversity of fitness vectors, i.e. how much of the fitness space is covered can be used. Our unified prioritization algorithm uses a combination of these two to lean towards data with strong fitness vectors while ensuring that the full criteria set is represented.

6. SCIENCE ALERT

Prioritization can be used for more than just data downlink decisions. It can also be used to initiate opportunistic science activities, or a "stop and call home" feature, if an extraordinary rock has been found.

Targets of high science value can be identified for additional instrument measurements. Prioritization that calls for opportunistic science is a wasted capability without a method of re-sequencing the rover to obtain the additional scientific observations requested. This ability for real-time opportunistic science requires integrating the prioritization module with the onboard planning and scheduling system.

The capability to identify and react to science events that were not initially scheduled is referred to as a science alert. A science alert involves identification of a science opportunity through data analysis and the modification of the rover activities to react to the opportunity through a planning and scheduling module. There is a spectrum of possible reactions to detection of a science opportunity. In this work we describe two of the reactions that have been implemented – the stop-and-call-home science alert and a request to take an additional science measurement.

As a rover is moving across a region, it is possible that unanticipated scientifically interesting targets may be encountered. The OASIS system identifies such targets based on gathered data. Scientists have designated that certain features or rocks with specific properties are extremely important. If such a rock is identified it is a valuable discovery, however it is imperative that false alarms be minimized. OASIS uses a model that considers how likely it is that the measured signal truly represents detection of the target signal based on the similarity of the two signals, as well as estimated measurement noise. This information is then combined with an importance rating for the target to identify signals that merit stopping further rover travel until a communication opportunity, i.e. stop and call home.

A second form of reaction is to request that additional data be taken on a target that appears to be of interest. In OASIS, this reaction may occur as a result of a target that is particularly unusual or novel. In this case, the data analysis submits a request to the OASIS planner to add a new activity to the rover's schedule. The new request may be for a new image or spectrometer read. Once the data analysis software has identified a set of new science targets, these targets are passed to onboard planning and scheduling software that can dynamically modify the current rover plan in order to collect the new science data.

This component takes as input the new set of science requests, the current rover command sequence (or plan), and a model of rover operations and constraints. It then evaluates what new science tasks could be added to the current plan while ensuring other critical activities are preserved and no operation or resource constraints are violated.

Planning and scheduling capabilities are provided in OASIS by the Continuous Activity Scheduling, Planning and Re-Planning (CASPER) system [9,10]. CASPER employs a continuous planning technique where the planner continually evaluates the current plan and modifies it when necessary based on new goal, state and resource information. Rather than consider planning a batch process, where planning is performed once for a certain time period and set of goals, the planner has a current goal set, a current rover state, and state projections into the future for that plan. Thus when a science alert request is received the plan is incrementally updated to accommodate, if possible, the new goal. The planner is responsible for maintaining a plan that will ensure the rover has sufficient resources to maintain its health and complete critical goals. After incrementally adjusting the plan in response to current resource levels, new goals and original goals, the rover activity sequence is updated enabling the rover to execute the new goals.

Since science alerts may involve several different levels of reaction, OASIS has been designed to support a spectrum of reactions. The most basic reaction is to adjust the rover plan so that the flagged data is immediately sent back to Earth for further analysis and the rover holds at the current position, delaying other non-critical tasks. This and the collection of additional data at a site have both been implemented. Future reactions include having the rover alter its path to get closer to objects of interest before taking additional measurements and/or scheduling a close contact measurement (such as with a microscopic imager). These operations would provide new data that could not be obtained through image analysis alone. The level of reaction allowed during mission operations will be determined by the constraints and goals of the rover mission and may likely vary over the course of a mission.

7. Conclusions and Future Work

Onboard, autonomous, science analysis systems are currently in development and can be a useful tool in maximizing the science return from a mission [11, 2].

As rovers are required to travel longer distances between science sites, the importance of traverse science increases. In order to allow the engineering team the freedom to rapidly move the rover from site to site, and still gain some science information from the traverse, a certain level of autonomy in the gathering of data will be required. OASIS is a system that is uniquely suited to this pursuit of traverse science.

Two important new capabilities have been added to the assist the scientist in extracting useful information from the many images taken during a long rover traverse: a unified prioritization algorithm that can output a prioritized list of images for downlink, and a science alert that can identify science opportunities and direct the rover to collect additional data.

Future Work

In future work, both the data analysis methods and the planning and scheduling capability of OASIS will be expanded. Extensive testing and validation are required in the near term as well. Perhaps the most exciting ground test validation will occur in early 2004 when OASIS will be given the opportunity to analyze MER images.

Currently, OASIS only extracts texture, albedo and shape features from the rocks that it finds. With the addition of range data and coordinate data, OASIS will also be able to provide rock size (and thus rock size distribution for all of the rocks that it finds) and rock location in the summary table that it downlinks at the end of the day. In the upcoming year, OASIS will also be able to extract color information from color images.

Further image analysis including estimation of soil and atmospheric properties are planned. In addition, integration of data from other instruments, such as a point spectrometer, can provide valuable information to the onboard system.

In the area of planning and scheduling an emphasis will be on broadening the range of possible reactions to new science opportunities. In the near term this will include modification of the initial rover path to acquire higher quality data on a target of interest.

8. Acknowledgements

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

We would like to thank our sponsors, the Intelligent Systems Program of CICT, the Interplanetary Network Directorate Program, and the AISRP. We would also like to acknowledge the hard work of the rest of the current OASIS team: Ben Bornstein of JPL, Marty Gilmore of Wesleyan University, and James Dohm of the University of Arizona. We also, in part, owe our current novelty detection prioritization algorithm to the work of Dennis De Coste at JPL.

This work began as a result of discussions between Eric Mjolsness (JPL), Steve Saunders (JPL) and Ray Arvidson (Washington University) in 1998.

Finally, our work could not proceed without the support and advice of scientists outside of the team, in particular: Matt Golombek, Albert Haldemann, Leslie Tamppari and Frank Palluconi.

References

[1] Andrew Mishkin, Jack Morrison, Tam Nguyen, Henry Stone, Brian Cooper, Brian Wilcox. Experiences with Operations and Autonomy of the Mars Pathfinder Microrover. *Proceedings of the 1998 IEEE Aerospace Conference*. Aspen, CO, March 1998.

[2] Rebecca Castano, Robert Anderson, Tara Estlin, Dennis Decoste, Forest Fisher, Daniel Gaines, Dominic Mazzoni, and Michele Judd, "Rover Traverse Science for Increased Mission Science Return," *Proceedings of the 2003 IEEE Aerospace Conference.* Big Sky, Montana, March 2003.

[3] Tara Estlin, Rebecca Castano, Robert Anderson, Dan Gaines, Forest Fisher and Michele Judd, "Learning and Planning for Mars Rover Science," *IJCAI 2003 workshop notes on Issues in Designing Physical Agents for Dynamic Real-Time Environments: World modeling, planning, learning, and communicating*, Acapulco, Mexico, August, 2003.

[4] Victoria Gor, Rebecca Castano, R. Manduchi, Robert C. Anderson, and Eric Mjolsness, "Autonomous Rock Detection for Mars Terrain", *Space 2001, American Institute of Aeronautics and Astronautics*, Aug. 2001.

[5] Rebecca Castano, T. Mann, and Eric Mjolsness, "Texture analysis for Mars rover images," *Proc.*

Applications of Digital Image Processing XXII, SPIE Vol. 3808, July, 1999.

[6] Martha Gilmore, Rebecca Castano, T. Mann, Robert C. Anderson, Eric Mjolsness, R. Manduchi, and R. S. Saunders, "Strategies for autonomous rovers at Mars", in *J. of Geophysical Res*., Vol. 105, No. E12, Dec. 2000 pp. 29223-29237.

[7] Justin Fox, Rebecca Castano, and Robert C. Anderson, "Onboard autonomous rock shape analysis for Mars rovers," *IEEE Aerospace Conference*, Big Sky, Montana, Mar. 2002.

[8] P. Smyth, "Clustering using Monte-Carlo cross-validation", *Proc. 2nd Int. Conf. on Knowledge Discovery and Data Mining*, AAAI Press, 1996.

[9] Tara Estlin, Forest Fisher, Daniel Gaines, Caroline Chouinard, Steve Schaffer, and Issa Nesnas. Continuous Planning and Execution for a Mars Rover. *Proceedings of the Third International NASA Workshop on Planning and Scheduling for Space*. Houston, TX, Oct 2002.

[10] Steve Chien, Russell Knight, Andre Stechert, Rob Sherwood, and Gregg Rabideau. Using Iterative Repair to Improve the Responsiveness of Planning and Scheduling. *Proceedings of the Fifth International Conference on Artificial Intelligence Planning and Scheduling*, Breckenridge, CO, April 2000.

[11] Virginia C. Gulick, R. L. Morris, M. A. Ruzon, and Ted L. Roush, "Autonomous image analysis during the 1999 Marsrokhod rover field test," J. Geophysical Research, Vol. 106, No. E4, pg 7745-7764, 2001.

BIOGRAPHY

Dr. Rebecca Castaño*, Supervisor, Machine Learning Systems (MLS) group at JPL, OASIS Technical team lead. She received her Ph.D. in Electrical Engineering from the University of Illinois with her dissertion in the area of computer vision. Dr. Castaño has been advancing the state of the art in onboard science analysis methods for the past five years and has been lead author on numerous publications in this field. From 1999 - 2001, Dr. Castano served as the application lead for phenomenological computational field geology efforts in the MLS Group. She is currently the technology lead for science data processing for the Autonomous Sciencecraft Experiment on the New Millenium Program's ST6 project. Dr. Castaño is also the Team Lead of the Onboard Autonomous Science Investigation System (OASIS) project and the subtask lead for Data Analysis on the automated Multi-rover Integrated Science Understanding System (MISUS). Her research interests include machine learning, computer vision and pattern recognition.*

Michele Judd, P.E.*, OASIS Research Manager and senior technical staff member of JPL's Earth and Space Sciences Division. Ms. Judd received her B.S. in Petroleum Engineering from Stanford University and her M.L.A. in Organizational Development from Southern Methodist University. She currently manages two additional research tasks in earthquake modeling and simulation, "Numerical Simulations for Active Tectonic Processes: Increasing Interoperability and Performance" and "Complexity Computational Environments: Data Assimilation SERVO Grid." A nationally certified Quality Manager and licensed professional engineer, Ms. Judd has over 15 years of experience in project management, both within JPL and for Mobil Oil.*

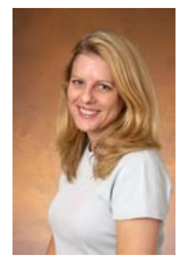

Dr. Tara Estlin, *OASIS Planning and Scheduling team lead, senior member technical staff. Dr. Estlin is a member of the Artificial Intelligence Group at the Jet Propulsion Laboratory where she performs research and development of planning and scheduling systems for rover automation and multi-rover coordination. Dr. Estlin is currently the PI of the "An Onboard Scientist for Multi-Rover Scientific Exploration" and the "Integrated Resource and Path Planning" projects. These efforts are developing capabilities for onboard rover-command generation, resource planning and scheduling, and data analysis for single and multiple rovers. She is also a member of the "Robotic Autonomy Architecture" project and is a past member of the Long-Range Science Rover (LRSR) project, which developed the Rocky 7 and Rocky 8 rovers. Dr. Estlin has been at JPL since 1998. She received a B.S. in computer science in 1992 fromTulane University, an M.S. in computer science in 1994 and a Ph.D. in computer science in 1997, both from the University of Texas at Austin.*

Dr. Robert C. Anderson*, OASIS Science team lead, senior member technical staff. Dr. Anderson attended Old Dominion University in Norfolk, Virginia, where he received his Bachelor of Science degree in geology in 1979. In 1985, he received a Master of Science from Old Dominion University in geology with an emphasis on structural geology and mapping tectonic features surrounding the Tharsis region of Mars. In 1995, he received a Doctor of Philosophy from the University of Pittsburgh in geology with an emphasis on visible and near infra-red remote* *sensing. His Ph.D. research was centering on mapping Quaternary surfaces and soils around the Whipple Mountains of southwestern Arizona. Dr. Anderson worked on the successful Mars Pathfinder Project as science support for the Mineralogy and Geochemistry Science Operations Group. Currently Dr. Anderson is the Investigation Scientist for the Rock Abrasion Tool (RAT) and science support for Mission Operations on the Mars 03 mission. He works closely with the FIDO rover team and is presently the Science Team lead on the Onboard Autonomous Science Investigation System (OASIS) project. Dr. Anderson's research is centered on unraveling the geologic history of Mars.*

***Lucas Scharenbroich** is currently a student in the School of Information and Computer Science graduate program at the University of California - Irvine. He is also an Associate Member of the Machine Learning Systems Group at JPL. He received undergraduate degrees in Computer Science (B.S.) and Electrical Engineering (B.S.E.E.) from the University of Minnesota, Duluth. His work at JPL has focused on applying unsupervised learning methods to space-based and bioinformatic problems.*

***Lin Song** is an associate member of the technical staff, Flight Software Systems and Technology Group at the Jet Propulsion Laboratory. She received a dual B.S. degree in Computer Science and Physics from Rensselaer Polytechnic Institute (RPI) in 2001. Ms. Song has worked on several projects in the applied math area since she came to JPL. She has developed a gravitational gradient model based on the spherical harmonics formalism to study gravity gradiometer measurements, implemented an inversion algorithm for optimal estimation of gravitational gradient parameters with measurement data. Currently Ms. Song is a member of the "Onboard Autonomous Science Investigation System" (OASIS) working on algorithm development. She is also working in the areas of software architectures, specifically explore and analyze the behaviors of a software system design using the SPIN model checker.*

***Dr. Daniel Gaines** is a senior member of the Artificial Intelligence Group at JPL. His research interests are in integrated planning and execution and in machine learning to improve planning. Dr. Gaines received a Ph.D. in Computer Science from the University of Illinois at Urbana-Champaign. Before coming to JPL, Dr. Gaines was an Assistant Professor in Computer Science at Vanderbilt University. His work at JPL is primarily focuses on planning and execution techniques for planetary exploration rovers.*

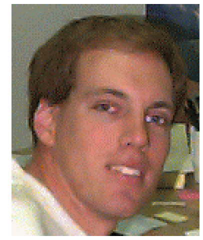

***Forest Fisher** is a senior member of the technical staff, Artificial Intelligence Group of JPL, California Institute of Technology. He holds a B.S. in Computer Science from the University of Texas at Austin, and is currently completing a M.S. in Computer Science at the University of Southern California. Mr. Fisher performs research and development of automated planning, scheduling and execution systems for autonomous rover operations and ground communication station automation. He is the PI of the "Unified Planning and Execution" project. He is also a member of the "Onboard Autonomous Science Investigation System" (OASIS) and the "Multi-Rover Integrated Science Understanding System" (MISUS) projects.*

***Dominic Mazzoni**, OASIS Research Programmer, associate staff, Machine Learning Systems. Mr. Mazzoni received his Bachelor of Science degree in Mathematics from Harvey Mudd College in Claremont, CA in 1999 and expects to receive his Master's Degree in Computer Science from Carnegie Mellon University in Pittsburgh, PA in early 2003. His Master's work focused on Music Information Retrieval, resulting in a working system for retrieving songs from a database of music using hummed queries. His work at JPL has focused on extending Machine Learning algorithms to work on massive implicitly-defined data sets.*

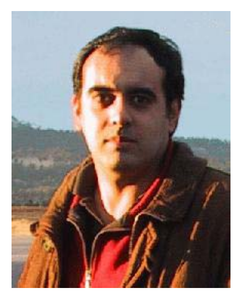

***Dr. Andres Castaño** received his B.S. and M.E.E. degrees from the University of Los Andes, Bogota, Colombia, and his M.S. and Ph.D. degrees from the University of Illinois at Urbana-Champaign, all in Electrical Engineering. At Urbana-Champaign he worked at the Beckman Institute in visual servo-control of manipulators and omni-directional cameras. In 1998 he joined the Information Sciences Institute at the University of Southern California where he was the lead roboticist of the Conro project on reconfigurable robots. In 2000 he joined the Machine Vision Group at the Jet Propulsion Laboratory where he currently works in feature detection and recognition. His research interests are robotics, computer vision and graph theory.*